Dual-Frequency Multi-Constellation (DFMC) SBAS Network Optimization Using Genetic Algorithms

Sam Pullen, Sukrut Oak, Sherman Lo, Juan Blanch, and Todd Walter
Stanford University
Mark Crews, Charles Frey, and Brice Barrett
Lockheed Martin

Abstract

This paper revisits the task of selecting optimal locations and numbers of Wide-Area Differential GNSS reference station (WRS) networks such as those provided by Satellite-based Augmentation Systems (SBAS). In the early days of the U.S. Wide Area Augmentation System (WAAS), a genetic algorithm (GA) was developed that selected the optimal combination of reference station locations from a list of potential locations using an early version of a service volume model (SVM) to assess the performance of each candidate set of locations. This approach generated relatively small and sparse reference station networks that proved to be inadequate to mitigate anomalous ionospheric spatial decorrelation events observed since 2000. Now that WAAS and other SBAS are being modernized to support multiple frequencies and constellations, the need for dense WRS networks is much less, and significant savings can be obtained by reducing their number.

This paper updates the original approach to this problem with a new decision variable, a new performance metric to be optimized, and updated GA operators using the Matlab Global Optimization Toolbox. The updated GA is tested and demonstrated with a simple performance metric that allows the optimal solution to be determined by inspection and compared with the GA result. Combining the GA with MAAST simulations of WAAS performance is conceptually the same, but the optimal solution is unknown, and the GA takes far longer to reach convergence due to the need to evaluate each of 30 solutions in the GA population at every GA generation. Preliminary results show that, as expected, less than half of the current 38 WRS are needed for dual-frequency operations, and the distribution of the retained WRSs is relatively uniform. This procedure can be extended to the design of any SBAS reference station configuration, which will be done in future work.

1. INTRODUCTION

When Satellite-based Augmentation Systems (SBAS) were originally developed in the early 1990s, the appropriate number and locations of the wide-area reference stations (WRSs) used to measure GNSS satellite ranges were not well known. One attempt to determine this via SBAS service volume modeling (SVM) tools and global optimization routines was reported in [Pullen, 1995b]. This approach used a genetic algorithm (GA) to evaluate potential SBAS network solutions and progress toward an optimal solution using the principles of biological evolution. This method showed that a surprisingly small number of WRSs (as few as 4 or 5) were needed to cover continent-sized areas in the U.S. and Europe. This solution was appropriate for the SVM used at the time (see [Pullen, 1995a]), which only considered small and gradual changes in ionospheric delay between satellite observations with ionospheric pierce points (IPPs) hundreds of kilometers apart. However, subsequent experience with severe ionospheric anomalies that affected SBAS in the early 2000s showed that protecting SBAS integrity during such events required many more WRS measurements made much closer together (see, e.g., [Walter, 2001]), which is why today's WAAS, EGNOS, and other continent-sized systems have several tens of WRSs separated by much less than 1000 kilometers.

Now, many years later, the prospect of removing the first-order effects of ionospheric delays from SBAS using dual-frequency and multi-constellation (DFMC) measurements presents us with an opportunity to revisit the SBAS network optimization problem with an eye to reducing the number of WRSs by more than 50% in existing systems and allowing new SBAS systems to be fielded much more inexpensively. An adaptation of the GA strategy used in [Pullen, 1995b] combined with the much more detailed SBAS SVM

included in the Stanford MAAST software [MAAST] that runs in Matlab is well-suited for this. MAAST evaluates SBAS performance based on the protection levels needed to meet the user integrity requirements for LPV and LPV-200 precision approaches, so it is a more reliable basis for performance optimization than the SVM derived in [Pullen, 1995a].

Figure 1 from [Pullen, 1995b] shows the elements of the GA optimization procedure developed in that paper. An initial population of 30 potential WRS solutions was generated randomly or by human selection to form "Generation 0", then the GA was allowed to "evolve" toward the optimal solution by using coverage results from the SVM in [Pullen, 1995a] (offline simulation-based integrity evaluations were used to augment the SVM but had little effect). The SVM results were provided to a value model that determined the quality of each WRS solution based on the projected availability of CAT I aircraft precision approach across the grid of user locations evaluated by the SVM. The GA operators (elitism, crossover, and mutation) were then applied to the population of solutions to (hopefully) "evolve" a set of improved solutions for the next generation. The procedure is repeated until convergence, which typically did not take more than several tens of generations running within 6 – 10 hours on a Sun workstation due to the relative simplicity of the SVM from 1995 compared to today's MAAST. The generation of solutions following convergence typically had many solutions that were the same – typically the one with the best performance (highest value) and thus the optimal result as determined by the GA.

Figure 2 is an updated version of Figure 1 showing the changes made to the GA procedure in [Pullen, 1995b]. The DFMC version of MAAST replaces the old SVM from 1995. In it, the standard 24-satellite GPS constellation almanac from [SPS, 2020] is used (this is essentially the same as the 24-satellite almanac defined in the WAAS MOPS [MOPS, 2020]), but no other constellations are included. Each MAAST simulation covers one day of repeatable GPS satellite geometries at 5-minute time intervals, giving a total of $12 \times 24 = 288$ epochs (and satellite geometries) per evaluation. Normally, MAAST is run with 1-to-2-minute intervals to sample the day more thoroughly, but this is infeasible for GA optimization because MAAST must be rerun for each of the 30 WRS location solutions in each generation of the solution population, and this must be repeated for every generation. With this small number of epochs per evaluation, MAAST can run and provide results in about 30 seconds for each solution (on a modern Windows 11 laptop), meaning that each generation takes a minimum of 15 minutes to process. In addition, the simplified objective value model from [Pullen, 1995b] has been replaced with one that separately evaluates performance and cost to provide a realistic tradeoff for the GA to consider. The GA operators used in [Pullen, 1995b] are retained with only minor changes but have been re-implemented in the Matlab Global Optimization Toolbox, which gives much more flexibility in modifying GA parameters.

Section 2 of this paper describes the WRS location problem statement in detail, including the definition of the solution vector to be optimized and the performance and cost functions used to determine the value of each candidate solution. Section 3 describes the GA operators used. Section 4 describes a simplified problem in which the value function based on MAAST is replaced by a simpler one for which the optimal solution can easily be determined by inspection and compared to the GA result for test and debugging purposes (since the lengthy run time required for MAAST evaluations is not needed). Section 5 shows the optimal solution found by applying the GA to the WAAS DFMC WRS localization problem. Section 6 shows the effects of manually modifying the optimal solution in a couple of ways. Section 7 summarizes this paper and describes the next steps for enhancing the methodology, reflecting more realistic operational constraints, and expanding it to additional augmentation systems.

2. WRS LOCATION OPTIMIZATION: PROBLEM STATEMENT

2.1 WRS Design Definition

While the optimization approach described in this paper can be applied to any SBAS network, The focus here is the modernization of the U.S. WAAS into a system that provides corrections and integrity information on two frequencies (L1 and L5) so that dual-frequency users can estimate and remove local ionospheric delays from their measurements. This should allow full or near-perfect (close to 100%) availability for LPV-200 flight operations with a subset of the 38 existing WRSs across the U.S., Canada, and Mexico shown in Figure 3. Each WRS location in this figure is assigned a number from 1 to 38 to identify it for the GA localization problem. The WRS design vector to be optimized is 38-element binary vector of ones and zeros, where a 1 represents retaining that WRS in the DFMC network, and a 0 represents omitting it. Thus, a vector of all 1s represents retaining all 38 existing WRSs, but the expectation is that many of the current WRSs will no longer be needed to achieve full LPV-200 availability. Note that any list of potential WRS locations can be used in this approach, and it would thus be easy to add new locations to the 38 existing ones (or change the locations entirely) if that were thought to be beneficial.

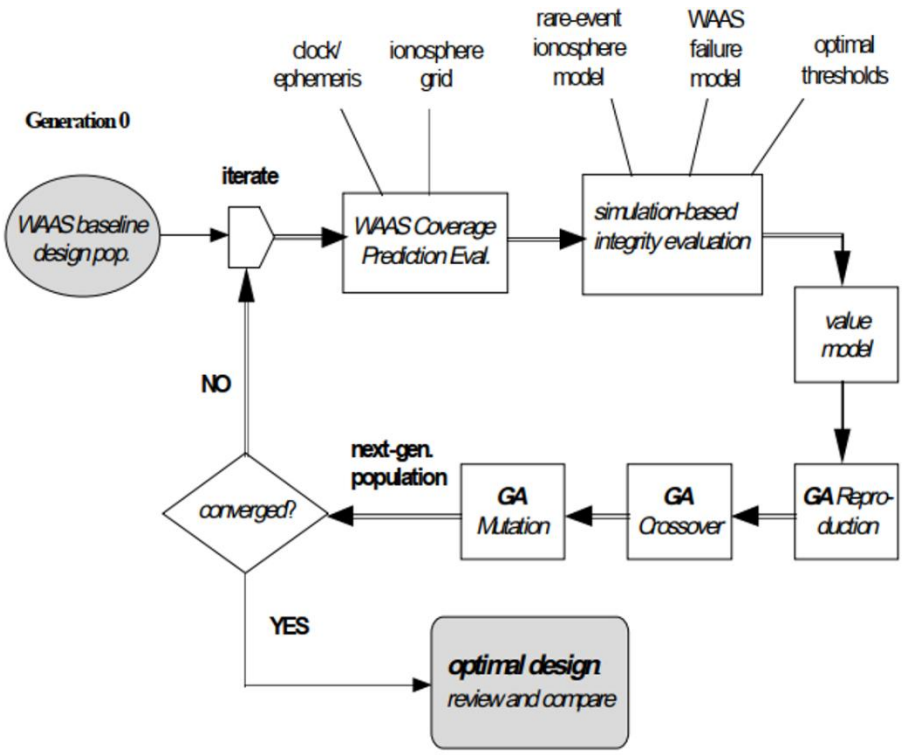

Figure 1: *SBAS Reference Station GA Optimization Procedure from [Pullen, 1995b]*

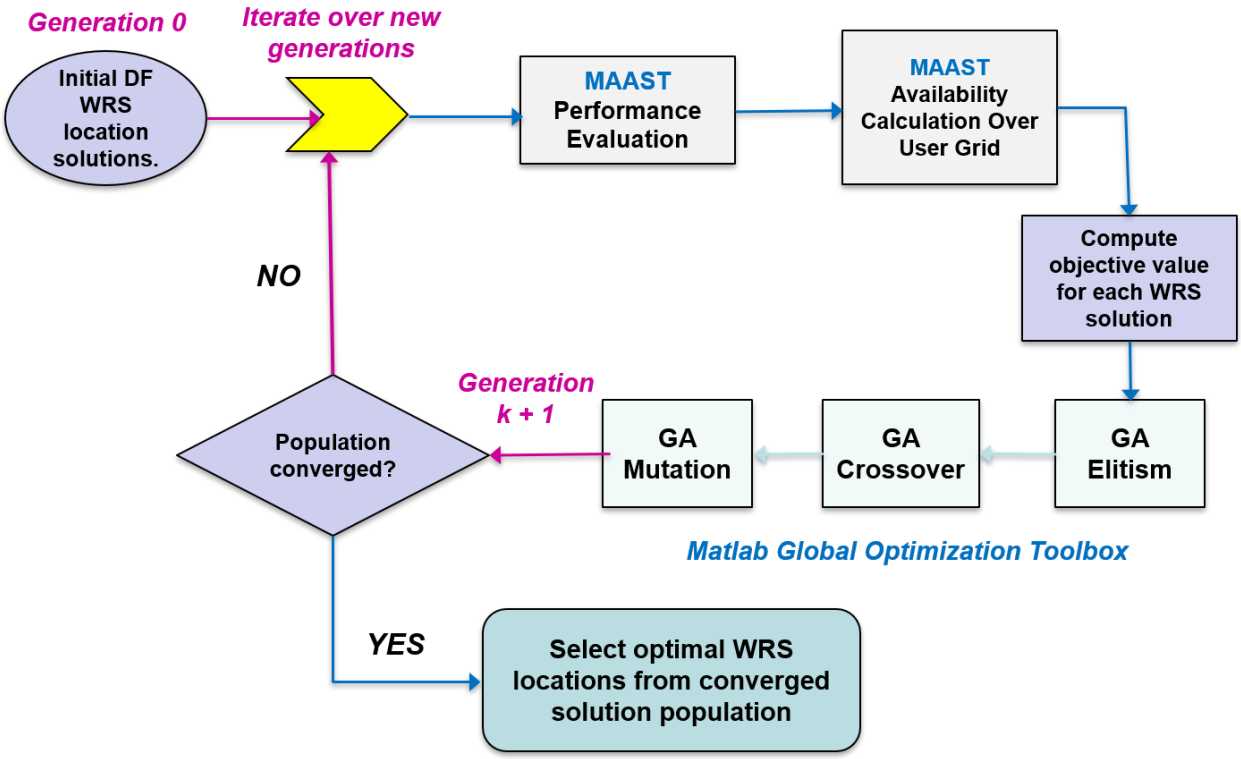

Figure 2: *Updated SBAS Reference Station GA Optimization Procedure*

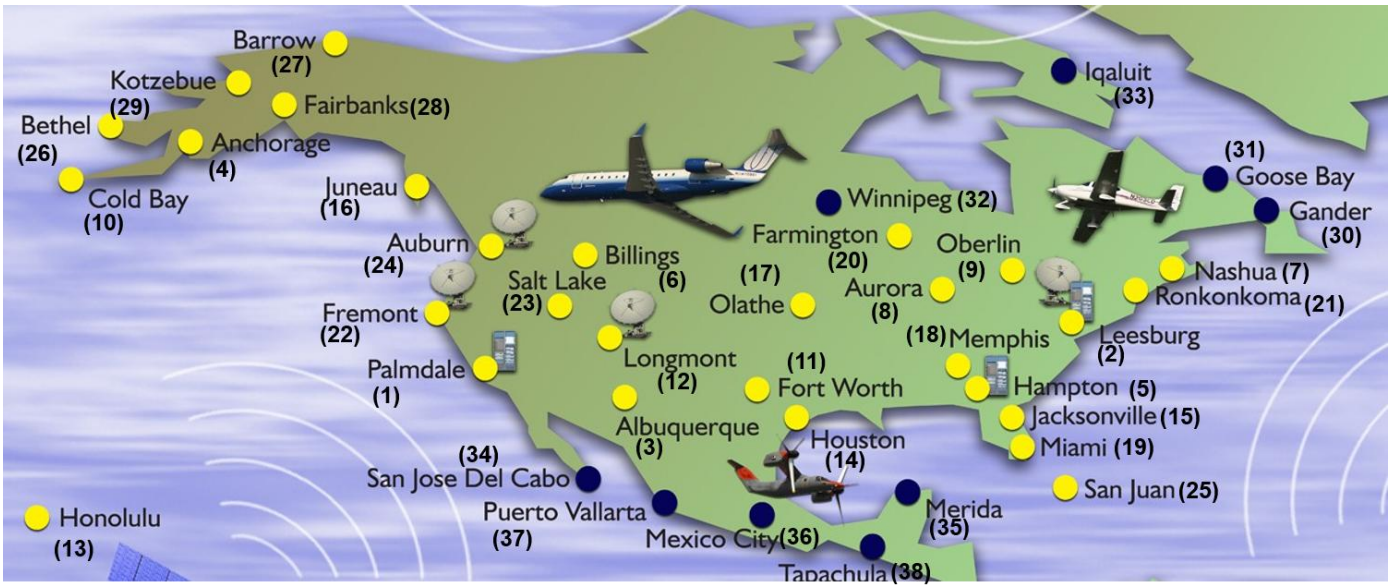

Figure 3: *Existing WAAS Reference Stations and Location Numbers Assigned for GA Optimization*

2.2 WAAS Life Cycle Cost Model

The actual cost structure for WAAS is not well-known, so a simplified approximation is used as part of the GA value model. The total life-cycle cost of the modernized DFMC WAAS ground system is assumed to be a fixed cost of US $5 billion plus $100 million for each WRS retained. Thus, the total cost of retaining all 38 existing WRSs would be $8.8 billion. Reducing the number of WRSs by half (to 19) would instead cost $6.9 billion, which is a significant savings that motivates the GA to remove as many stations as possible without unduly sacrificing performance. Note that this model implies that the optimized SBAS and WRS network will be implemented from scratch, as the economics of modifying the existing WAAS network to reflect the need for fewer WRSs under dual-frequency operation are quite different (e.g., the cost savings from removing existing stations are likely to be smaller than those from avoiding the need to install new stations).

2.3 WAAS Performance Model and Overall Value Calculation

The performance metric used in this model is based on the presumed economic benefit of the LPV-200 approach availability provided by MAAST across the 2.5 × 2.5-degree latitude and longitude grid of user locations from 25 to 55 degrees North latitude and 65 to 135 degrees West longitude, which covers all of the Conterminous U.S. (CONUS), southern Canada, northern Mexico, and nearby oceans. Perfect (100%) availability at all user locations translates into a maximum total user benefit of $10 billion over the DFMC WAAS life cycle. This is a very conservative assessment of WAAS user benefits over 30 years. The reason for selecting $10 billion as the maximum benefit is to keep it in the same range as the total cost. If benefits were allowed to reach or exceed $20 billion, they would dwarf the presumed costs and greatly reduce the significance of retaining or omitting WRS locations.

To more severely penalize LPV-200 availability below 99%, the multi-step function shown in Figure 4 is used to translate the LPV-200 availability result into a multiplier *B* that is applied to the maximum benefit of $10 billion. *B* is a maximum of 1.0 for perfect availability but falls off rapidly as availability drops toward or below 90%. Figure 4 is derived from the Matlab code shown below that computes *B* ('perf_value') from availability (performance', expressed as a percentile with a maximum of 100%):

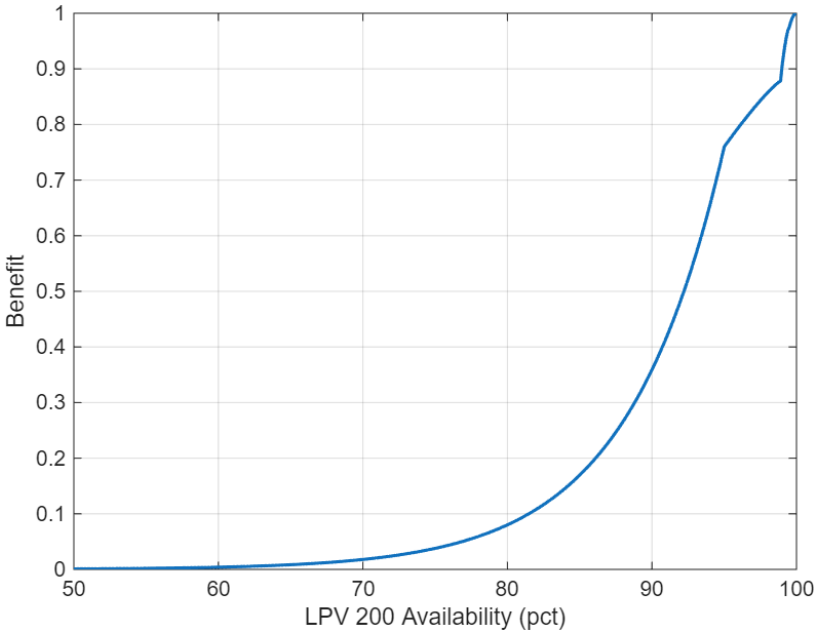

Figure 4: *Benefit Multiplier (B) Based on LPV-200 User Availability*

```
if performance >= 100
    perf_value = 1.0;
elseif performance >= 99.5
    perf_value = 1.0 - 0.05 * (100 - performance) / 0.5;      % Slow decrease from 100% to 99.5%

  elseif performance >= 99.0
      perf_value = 0.95 - 0.10 * (99.5 - performance) / 0.5;    % Moderate decrease from 99.5% to 99%
  elseif performance >= 95.0
      perf_value = 0.85 - 0.45 * (99 - performance) / 4.0;       % Faster decrease from 99% to 95%
  else
      perf_value = 0.4 * exp(-0.08 * (95 - performance));       % Rapid decrease below 95% using exponential decay
  end
```

Each location in the user grid defined above has its performance value computed from its LPV-200 availability. To generate a composite performance value over the entire grid, a weighted function is applied that more severely penalizes individual locations with the lowest availability in the user grid. This motivates the GA to retain enough WRS locations to avoid leaving parts of the user grid with much poorer than average availability. The following calculation is implemented on the availability results from each of N grid locations:

a) The location with the lowest availability among the N locations has its performance value multiplied by 10.
b) The 10 locations with the next-lowest availability have their performance values multiplied by 3.
c) The remaining $N - 11$ locations would have their performance values multiplied by 1 (unweighted).
d) To produce a composite result (B_{comp}) in the range of 0 to 1, the sum of the multiplied performance values derived from steps (a, b, c) are divided by the sum of the multipliers: $10 + (10 \times 3) + (N - 11) = (N + 29)$.

The final value function applied by the GA is then:

$$\text{Value} = (\$10 \text{ billion})\, B_{comp} \;-\; (\$5 \text{ billion}) \;-\; (0.1 \text{ billion})\, N_{WRS}$$

where N_{WRS} is the number of retained WAAS reference stations (out of 38).

3. GENETIC ALGORITHM (GA) DESIGN

A GA maintains not one but many candidate solution vectors in a population that evolves toward the optimal solution over many "generations" (iterations). As explained above, each solution vector is 38 elements long, corresponding to each of the 38 active WRSs in the U.S. Wide Area Augmentation System (WAAS) shown in Figure 3. Each "1" represents the presence of the corresponding WRS in the existing WAAS network, whereas a "0" represents its absence (in other words, that WRS location is not present in the candidate solution in question). The initial GA population of 30 candidate solutions or "members" of generation 0 includes 1 solution with all 38 existing WRSs (all "1s"), 4 other solutions with manually selected subsets of these 38 WRSs, and 25 random vectors where each element is equally likely to be a "1" or "0".

Variations of the three traditional GA operators are applied to go from each generation of member solutions to the next. The first, "elitism", simply selects the 5 members with the highest value (as described in Section 2) and reproduces them unchanged in the next generation. The second, "crossover", simulates the mixing of genes through sexual reproduction in biological evolution. The crossover operator randomly selects 25 pairs of members (parents') based on fitness value. Specifically, higher-valued members are more likely to be selected, and the 10 members with the lowest values are excluded. For each set of parents, a random point in the 38-element vector is selected, and their 'gene' values (0s and 1) are exchanged to generate two 'child' vectors that are members of the next generation. A simplified view of this process for 14-element members is shown in Figure 5, where Child 1 is created from Parent 1's elements to the left of the crossover point and Parent 2's elements to the right of it. Child 2 is generated analogously from the left-hand elements of Parent 2 and the right-hand elements of Parent 1. The third operator, "mutation", operates on the 30 candidate members of the new generation by randomly flipping each it (1 to 0 or 0 to 1) across the 30 members with a small probability of 0.01. GA "evolution" continues until the 5 population members with the highest values stop changing much ('convergence') or a maximum of 60 generations is reached, representing approximately 15 hours of GA run-time on a modern Windows 11 laptop that is weighted down by other simultaneously running processes.

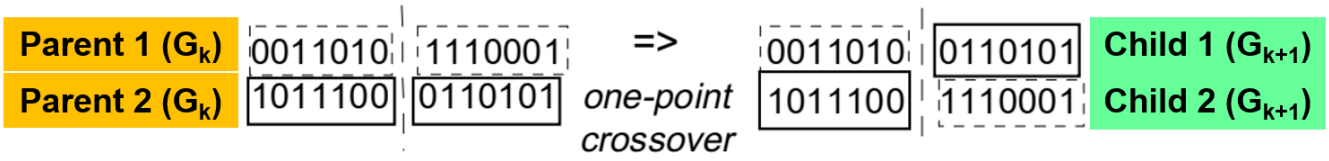

Figure 5: *One-Point Crossover Operator Applied within Genetic Algorithm*

4. GA TESTING ON A SIMPLIFIED PROBLEM

Because the MAAST evaluations of LPV-200 availability for DFMC WAAS take so long, the Matlab GA operators coded in Matlab are tested on the same problem with a much simpler performance model that can be computed quickly and analytically without any need for MAAST. The same 38-element solution vector, GA operators, and cost model are used, but the performance model is replaced with the following:

a) The availability of each solution starts at 100%.
b) Each omitted WRS location in an *odd* location (e.g., the 1st, 3rd, 5th, etc. locations in the solution vector) results in an availability reduction of 0.1%.
c) Each omitted WRS location in an *even* location (e.g., the 2nd, 4th, 6th, etc. locations in the solution vector) results in an availability reduction of 0.5%.

d) Furthermore, each omitted WRS location in a location divisible by 4 (e.g., the 4th, 8th, 12th, etc. locations in the solution vector) results in an additional availability reduction of 0.5%. Since these locations are also even numbers, the total reduction is 1.0%.

e) The resulting availability following these reductions is input into the multiplier function defined in Section 2.3 to derive the final benefit B_{comp}.

f) The value function shown in Section 2.3 is then used to derive the value of each solution.

Since the cost of each WRS is the same regardless of its location in the solution vector, it is easy to see that odd-numbered WRS locations should be removed before any even-numbered ones are. Further, only the numbers of retained odd, even, and divisible-by 4 are needed to evaluate each solution. Straightforward enumeration of the possible solutions produces the list of the top six solutions shown in Table 1, where the performance, cost, and value results are in billions of dollars. All six solutions retain all nine WRS locations divisible by 4, as removing any of these results in a severe availability penalty. The best solution (rank 1) removes all odd-numbers locations and retains all even-numbered ones, meaning that it represents a unique vector of 1s and 0s. The other solutions shown in Table 1 are not distinct because many different combinations of 1 or 2 odd locations retained or even locations removed exist. Note that, while the unique top-ranked solution has the highest value, the fourth-ranked solution has the lowest cost (because two even-numbered locations are removed) of the six solutions shown, while the sixth-ranked solution (which retains two odd-numbered locations) has the highest performance of those shown.

Table 1: *Optimal WRS Solutions for the Simplified Performance Model (Values in $ Billion)*

Rank	n_{odd}	$n_{div2not4}$	n_{div4}	Perf.	Cost	Value
1	0	10	9	8.60	6.90	1.70
2	0	9	9	8.46	6.80	1.66
3	1	10	9	8.63	7.00	1.63
4	0	8	9	8.31	6.70	1.61
5	1	9	9	8.49	6.90	1.59
6	2	10	9	8.65	7.10	1.55

When this value model is applied to the GA operators defined in Section 3, Figure 6 shows the resulting GA population evolution. The GA process runs in less than 30 seconds (due to no need for MAAST evaluations) and typically converges to the known optimal solution before the 60th GA population generation. Note that the Matlab GA implementation minimizes the objective function provided, thus what is shown is minimizing "cost minus benefit" such that the lowest negative value (– $1.7 Billion) is the best result. While the best solution achieves the optimal result, most solutions in the GA population are significantly worse, and the mean solution value stops improving around the 15th generation. For the best solution of the current GA generation to improve upon the previous generation, a new solution newly generated in the current generation by the "crossover" and/or "mutation" functions must be better than the five previously best solutions copied over into the current generation by the "elitism" operator.

In general, GA solution improvement does not happen at every generation but instead requires the right combination of factors to occur, which happens only sporadically. This is due to the random nature of the GA crossover and mutation operators. While all other aspects of the procedure (including MAAST evaluation of each candidate solution) are deterministic, this randomness means that different population evolutions will result from different initial "seeds" of the Matlab random number generator. For the simplified problem, running the GA procedure many times using different initial seeds showed that the known optimal solution is reached within 60 generations approximately 70% of the time. If left to run indefinitely, this randomness should theoretically result in eventually obtaining the optimal solution in all cases.

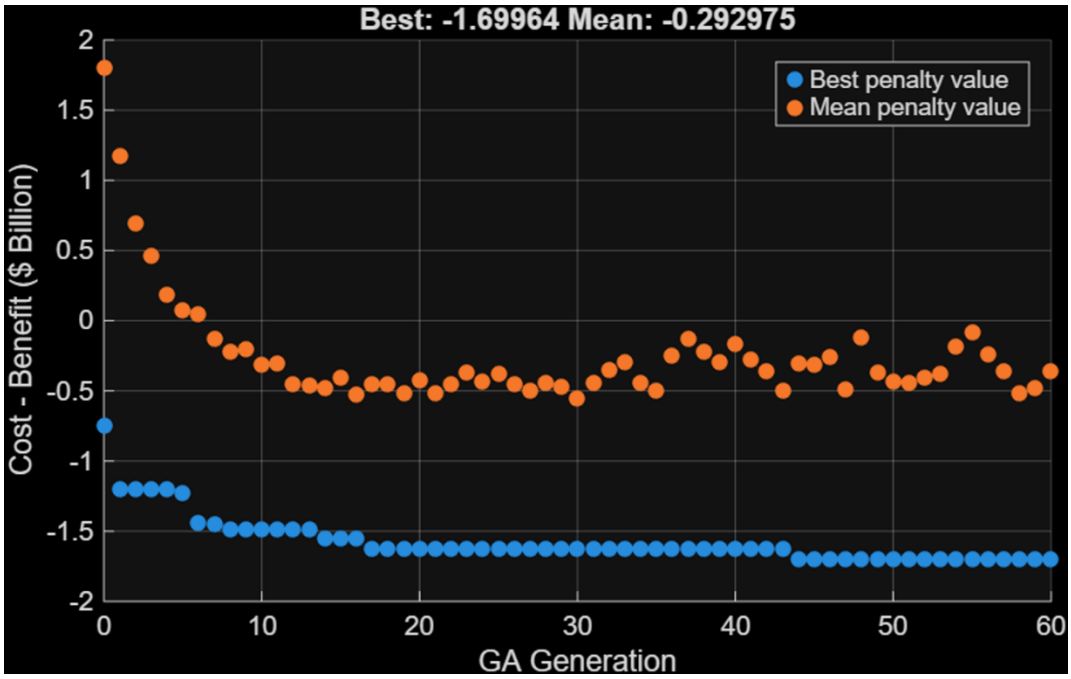

Figure 6: *GA Population Evolution and Objective Values for Simplified Problem*

5. GA IMPLEMENTATION WITH MAAST EVALUATION

Figures 7 and 8 show the WAAS LPV-200 approach availability determined by MAAST for both single-frequency (Figure 7) and dual-frequency (Figure 8) WAAS with all 38 existing WRSs shown in Figure 3 across the user grid evaluated by the GA. These results (and all that follow) are based on the standard 24-satellite GPS constellation from [SPS, 2020]. Comparing these two figures shows the advantages of the dual-frequency approach. While the existing WAAS network has been designed and refined over time to provide nearly 100% availability of LPV-200 in and around CONUS, southern Canada, and much of Alaska, the switch to dual-frequency provides perfect 100% availability at every location in the user grid (under the assumptions made here, without satellite or WRS outages). This strongly suggests that, as will be shown below, many fewer WRSs could provide nearly the same availability at lower cost.

The full 38-WRS configuration is one of the five manually-defined solutions included in the initial GA solution population in Generation 0 and has an objective value of $ 10 – 5 – 38(0.1) = $ 1.2 Billion. Availability results for two of the other initial solutions are shown in Figures 9 and 10. Figure 9 shows the result obtained by arbitrarily retaining all 19 even numbered WRSs using the station numbers in Figure 3. Retained WRS locations that lie within the user grid evaluated by the GA are shown on this plot, but others exist outside this grid and are not shown. This solution has an objective value of $ 9.6 – 5 – 19(0.1) = $ 2.7 Billion, which is much better than retaining all 38 WRSs but does not represent an optimal WRS distribution, as the eastern quadrant of the user grid is weakened by the absence of any WRSs on the east coast of CONUS other than the one near Washington, D.C.

The 14-WRS solution used in Figure 10 removes four of the 19 stations in Figure 9 and moves the Billings WRS to Boston (an odd-numbered location) to attempt to rectify this eastern availability loss. It is successful in doing so, as availability around and off the East Coast is much better, while the loss of availability near the West Coast is small and is well out in the ocean. This solution has an objective value of $ 9.7 – 5 – 14(0.1) = $ 3.3 Billion.

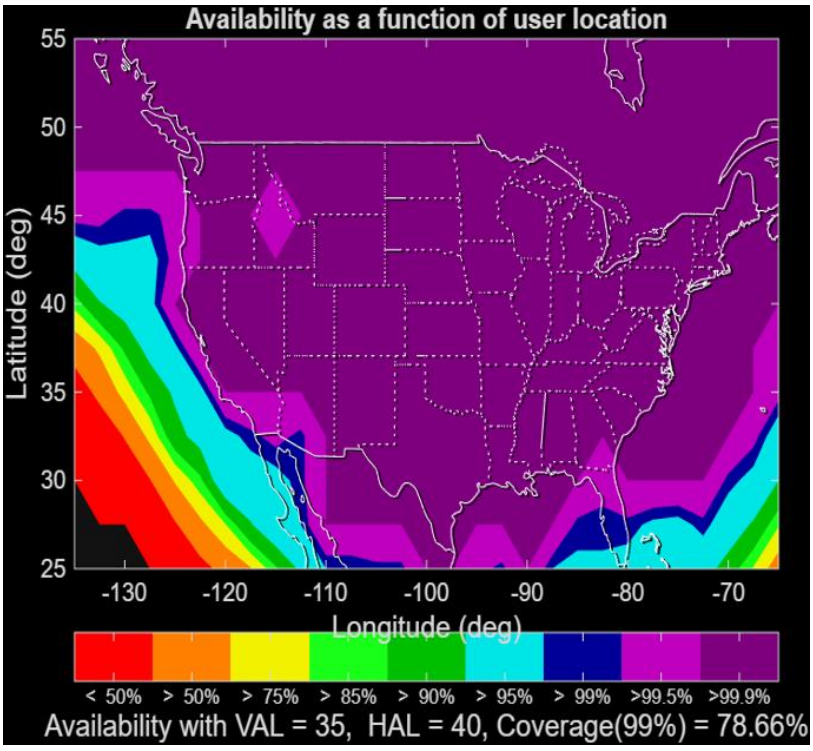

Figure 7: LPV-200 *Availability Results for Single-Frequency (L1-only) WAAS with 38 WRS Locations*

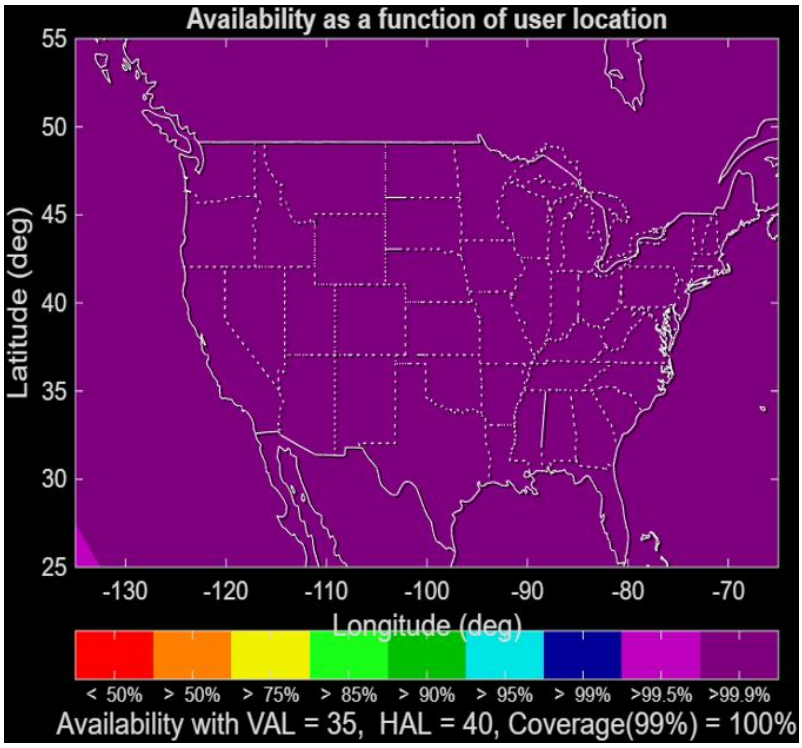

Figure 8: LPV-200 *Availability Results for Dual-Frequency (L1/L5) WAAS with 38 WRS Locations*

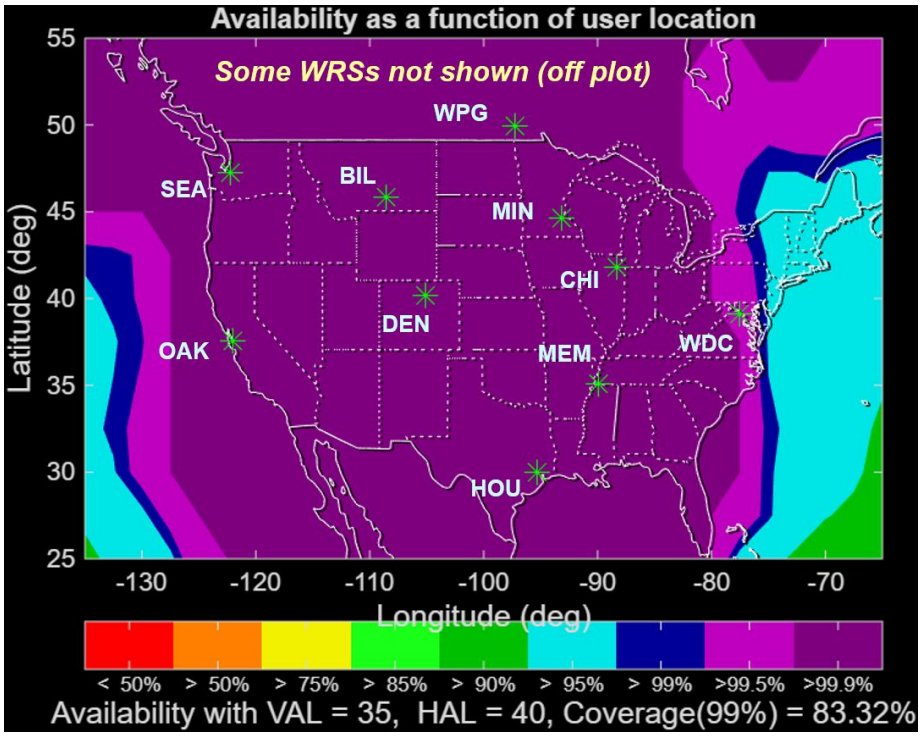

Figure 9: *LPV-200 Availability Results for Initial Solution with 19 Even-Numbered WRS Locations*

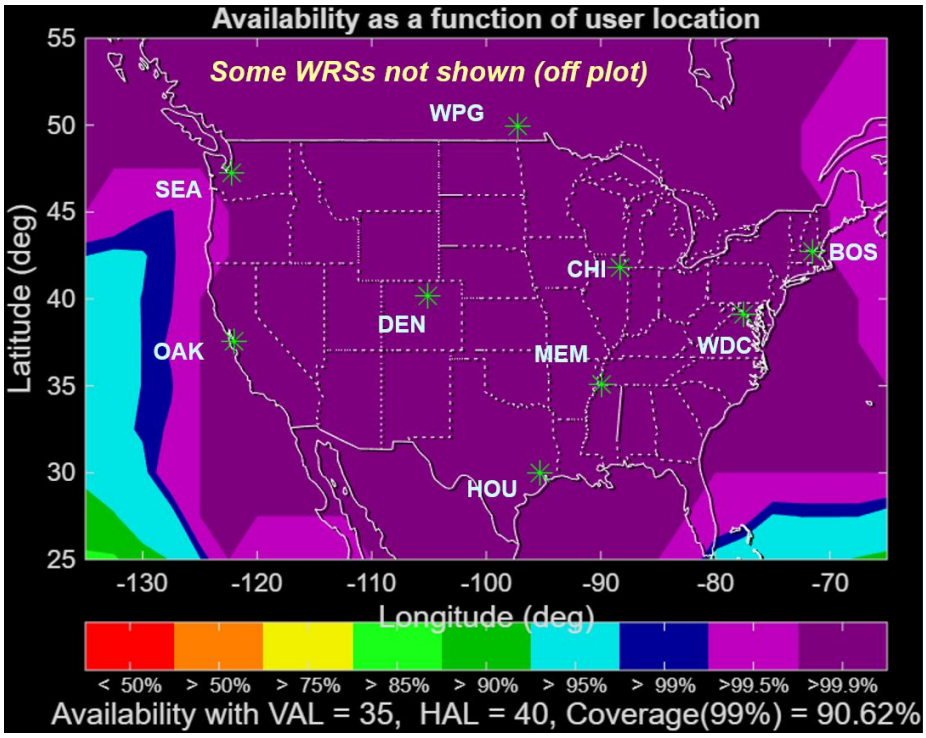

Figure 10: *LPV-200 Availability Results for Initial Solution with 14 of Even-Numbered WRS Locations and Billings replaced by Boston*

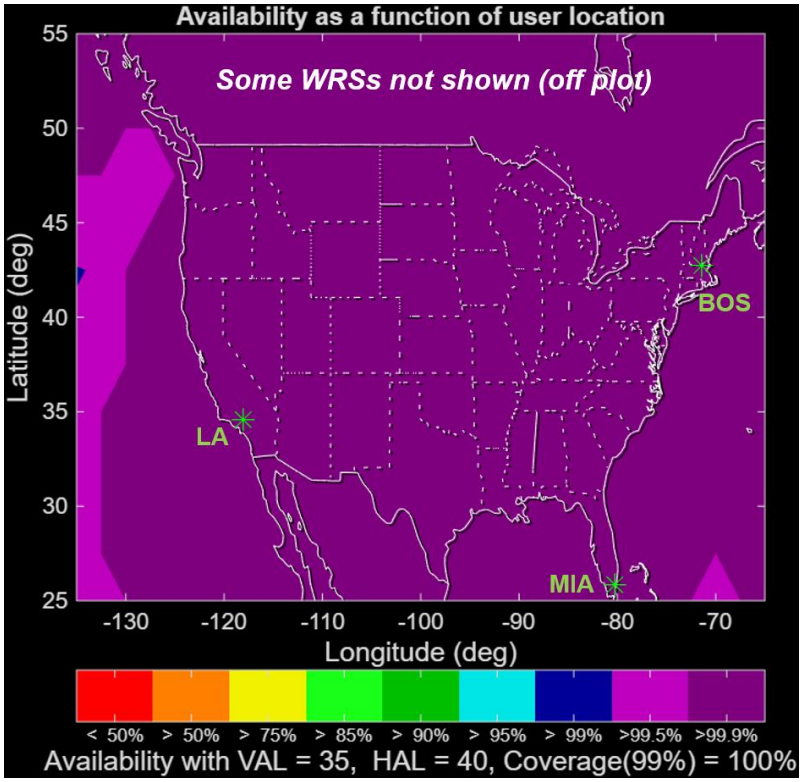

Figure 11: *LPV-200 Availability Results for Optimal GA Solution with 9 WRS Locations (over user grid evaluated by GA)*

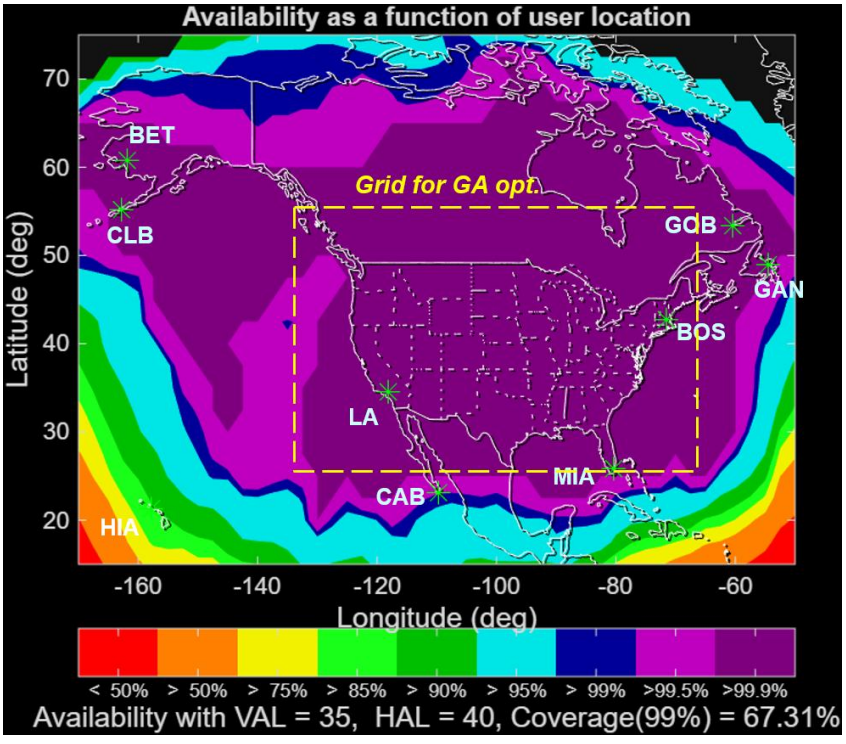

Figure 12: *LPV-200 Availability Results for Optimal GA Solution with 9 WRS Locations (over larger user grid)*

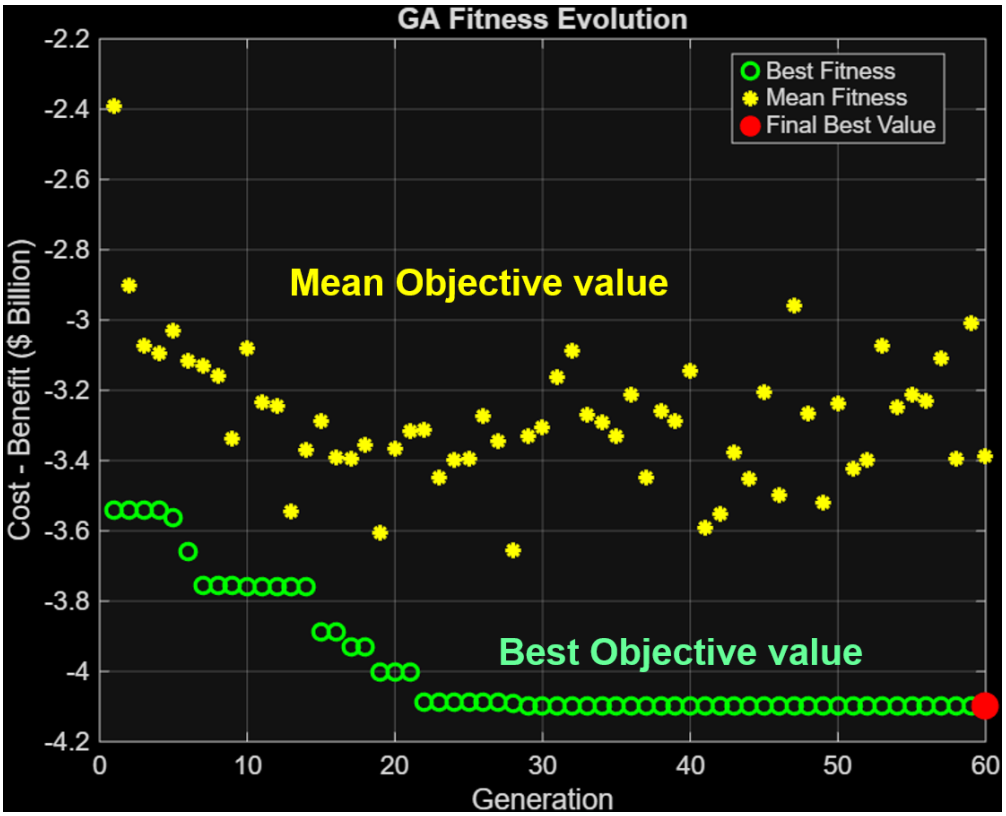

Figure 13: *GA Population Evolution and Objective Values for WRS Location Optimization*

Figures 11 and 12 show the best WRS solution obtained from multiple runs of the GA process with MAAST availability evaluation. It uses only 9 WRSs, of which only three are inside CONUS and are shown in Figure 11. Figure 12 shows the same result over a larger user grid extending from 15 to 75 degrees North latitude and 50 to 170 degrees West longitude, which includes all 9 WRSs and indicates the smaller zone used for GA optimization inside a white-hashed-line box. Despite the small number of WRSs, availability within the user grid evaluated by the GA in Figure 11 is nearly perfect and earns a performance value of just under $ 10 Billion ($ 9.998 Billion), giving an overall objective value of about $ 10.0 – 5 – 9(0.1) = $ 4.1 Billion. Figure 12 shows that availability close to 100% is maintained well outside this user region, with availability exceeding 99% (dark blue, light purple, and dark purple contours) over almost all North America excluding southern Mexico.

Figure 13 is analogous to Figure 6 and shows the evolution of mean and best performance (in terms of minimizing Cost – Benefit) within the GA population over 60 generations. This process took several hours on a 6-year-old Windows-11 PC with 32 GB of RAM that had few other processes running to maximize the RAM available to Matlab (some runs aborted earlier due to Matlab memory management issues). As in Figure 6 for the simplified problem, the mean objective value of the GA population stops improving after around the 20th generation, while the best objective value is lowered in streps (as new "best" solutions are generated by the GA) until the 29th generation, after which no further improvement is found. Examining the best five solutions in the final GA population shows that all of them use 9 WRSs, and all have no more than two different WRS locations compared to the solution in Figures 11 and 12 (for example, the 2nd-best solution replaces the Gander WRS with one in Cleveland). All five of these solutions have objective values exceeding $ 4.08 Billion, indicating that the peak of the objective value curve is relatively flat.

6. MANUAL MODIFICATIONS OF THE OPTIMAL SOLUTION

For many reasons, obtaining the optimal or near-optimal WRS configuration from the GA is only one step in the process of SBAS system design. For one thing, the performance and cost functions defined in Section 2 are only crude approximations of the desired optimal behavior and can and should be refined by humans once the GA finds the best solutions for these functions. Second, as

explained above, the GA procedure is highly random and may not find the best possible solution in a practical number of generations. Third, and perhaps most importantly, the MAAST evaluation of WAAS performance is limited to the condition where all 24 primary-slot GPS satellites and all WRSs are operating normally. This is not the case in practice, thus the operational SBAS network should make allowances for at least one and perhaps two WRSs having outages at any given time.

Figure 12 shows that the optimal set of WRSs identified by the GA includes two pairs of WRSs relatively close to one another: Gander and Goose Bay in eastern Canada and Bethel and Cold Bay in Alaska. It is natural to wonder if this is truly optimal or if one WRS in each pair might be redundant. Figure 14 shows the MAAST-determined LPV-200 availability for a 7-WRS network in which the Goose Bay and Cold Bay stations have been removed. Note that availability near the edges of the GA-evaluated user grid (and well beyond) is much worse than in the optimal solution, resulting in a loss of $ 1.2 Billion in performance that greatly exceeds the $ 0.2 Billion cost saved from having two fewer WRSs. The reason for this is that each GPS satellite needs to be visible to multiple WRSs simultaneously for clock/ephemeris corrections with low User Differential Range Errors (UDREs) to be provided by WAAS. This creates a lower bound in the number of WRSs over North America needed to provide nearly 100% LPV-200 availability.

Figure 15 goes in the other direction and attempts to address the potential lack of robustness in the optimal solution due to its assumption that all 9 WRSs are always operational. In it, three WRS locations in Seattle, Billings, and Cleveland are added to the optimal set of 9 WRSs. Comparing Figure 15 to Figure 12 shows that adding these three WRSs improves availability in many places (for example, many locations with availability between 99.5% and 99.9% in Figure 12 now have availability exceeding 99.9%), but this improvement matters little in the performance model defined in Section 2.3. The real benefit of the added WRSs is robustness to individual WRS outages, which MAAST will be used to evaluate in future work.

7. SUMMARY AND FUTURE WORK

This paper describes the elements of a GA optimization procedure to determine optimal WAAS reference station locations based on an earlier attempt in [Pullen, 1995b] along with an updated evaluation tool (DFMC MAAST) and expanded value function that combines performance (in terms of LPV-200 user availability) and cost (which is reduced if fewer WRS locations are retained from the existing set of 38). The GA performs well on the simplified problem defined in Section 4 and finds multiple improved solutions when paired with the DFMC MAAST in Section 5. The resulting optimal set of 9 WRS locations (and the many minor variations of this set that have nearly the same objective value) is significantly smaller than the manually selected WRS solutions used to initialize the GA while still providing LPV-200 availability close to 100% over the grid of user locations considered by the GA. The GA combined with MAAST takes hours to run and has a significant random element that prevents the GA from being guaranteed to converge to the best solution in a reasonable time frame (60 to 100 generations), but it is a very useful tool in guiding human designers toward the numbers and locations of WRSs likely to be optimal in practice.

As explained in Section 6, even if the GA were known to converge to the optimal solution defined by the selected performance and cost functions, further investigation by human designers is needed. Refinements to these functions should be examined, including a more realistic assessment of the marginal cost of each added WRS and a modification to the function in Figure 4 to further penalize good-but-not-perfect user availability between 99% and 99.9%. More importantly, the MAAST evaluation should be extended to consider the worst performance with a single WRS outage. The simplest way to do this is to repeat the MAAST evaluation of each candidate solution N_{WRS} times with each of the selected WRS locations set to be unavailable. This would of course extend the GA processing time by a factor of 10 or more, but a partial countermeasure is possible by maintaining a list of candidate solutions that have already been evaluated so that MAAST can skip repeated evaluations of the same solution (as these will obtain the same result).

This WRS location optimization procedure can easily be extended to dual-frequency SBAS network design in other regions of the world, particularly those in equatorial regions that do not have SBAS coverage today due to the limitations of single-frequency SBAS in handling equatorial ionospheric behavior. It can further be extended to include other elements of augmentation, such as satellites and ground transmitters used to convey differential corrections, and to examine optimal configurations for Precise Point Positioning (PPP – see [Lai, 2026] for another approach to this) and Real Time Kinematic (RTK) networks.

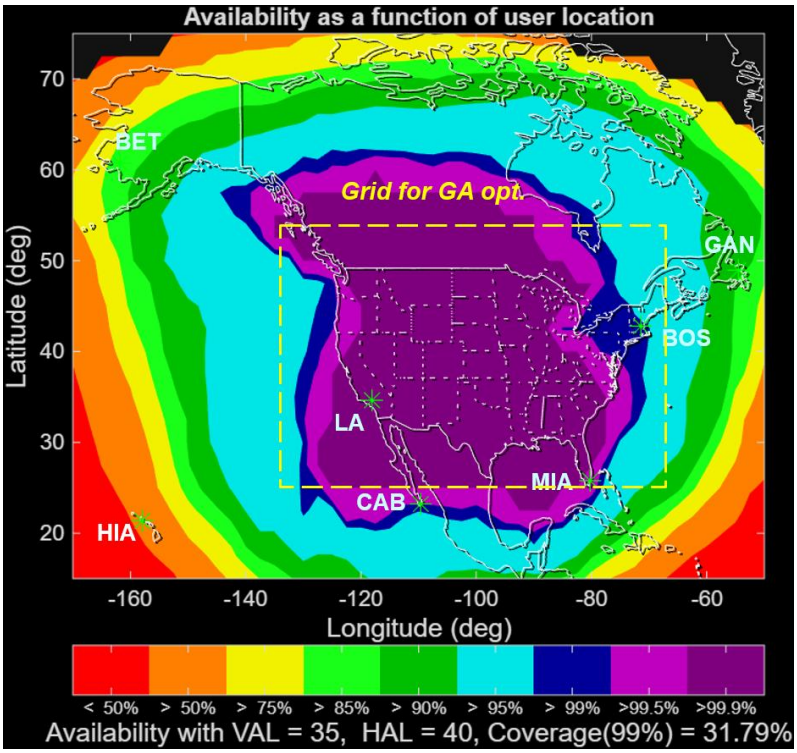

Figure 14: *LPV-200 Availability Results for 7-WRS Subset of the Optimal 9 WRS Locations*

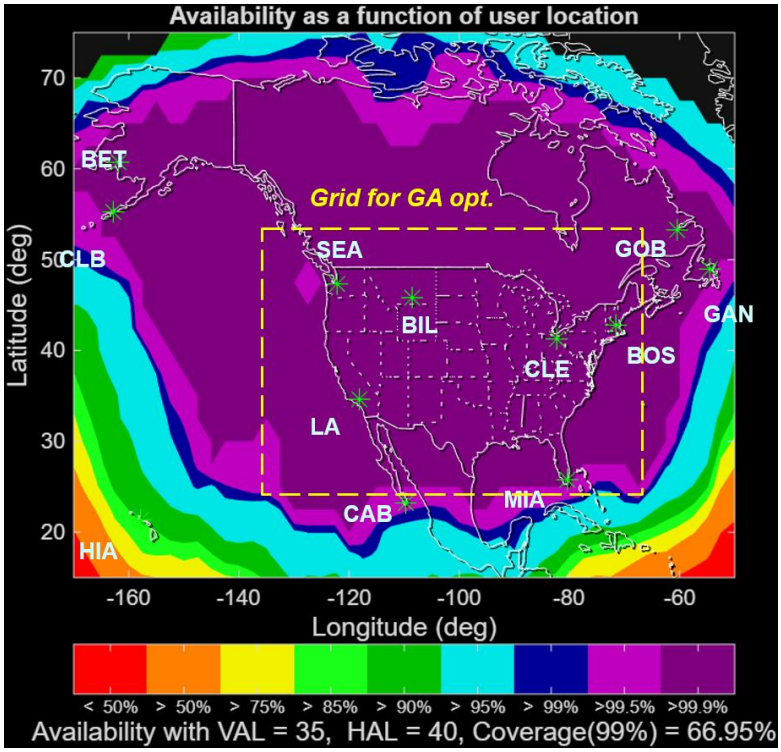

Figure 15: *LPV-200 Availability Results for 12-WRS Superset of the Optimal 9 WRS Locations*

8. REFERENCES

[Lai, 2026] Y.F. Lai, J. Blanch, and T. Walter, "Receiver Placement and Number Selection of Ground-Based Integrity Monitor via Convex Optimization," *Proceedings of ION GNSS+ 2026*, Orlando, FL, September 2026 (forthcoming).

[MAAST] Matlab Algorithm Availability Simulation Tool (MAAST), Stanford University GPS Laboratory, Dept. of Aeronautics and Astronautics. https://gps.stanford.edu/resources/software-tools/maast

[MOPS, 2020] *Minimum Operational Performance Standards (MOPS) for Global Positioning System/Satellite-Based Augmentation System Airborne Equipment*. RTCA SC-159, DO-229F, June 2020.

[Pullen, 1995a] S. Pullen, P. Enge, and B. Parkinson "A New Method for Coverage Prediction for the Wide Area Augmentation System (WAAS)," *Proceedings of 1995 Annual Meeting*, Colorado Springs, CO, June 1995. http://web.stanford.edu/group/scpnt/gpslab/pubs/papers/Pullen_IONAM_1995_waas_coverage.pdf

[Pullen, 1995b] S. Pullen, P. Enge, and B. Parkinson "Global Optimization of GPS Augmentation Architectures using Genetic Algorithms," *Proceedings of ION GPS '95*, Palm Springs, CA, September 1995. http://web.stanford.edu/group/scpnt/gpslab/pubs/papers/Pullen_IONGPS_1995_Global_Optimize.pdf

[Walter, 2001] T. Walter, A. Hansen, et al., "Robust Detection of Ionospheric Irregularities," *Navigation*, Vol. 48, No. 2, Summer 2001, pp. 89-100, DOI 10.1002/j.2161-4296.2001.tb00231.x,

[SPS, 2020] *GPS Standard Positioning Service (SPS) Performance Standard (GPS SPS PS)*, Washington DC, U.S. Dept. of Defense, 5th Edition, April 2020. https://www.gps.gov/technical/ps/2020-SPS-performance-standard.pdf